

AD-A191 391

PROBLEMS IN INTEGRATED NETWORKS(U) COLUMBIA UNIV NEW
YORK DEPT OF ELECTRICAL ENGINEERING M SCHWARTZ ET AL.
15 FEB 88 N00014-85-K-0371

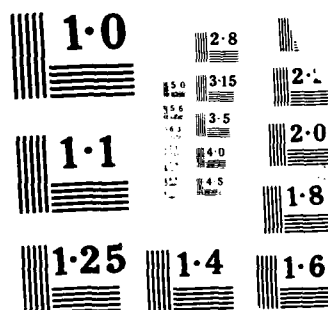
1/1

UNCLASSIFIED

F/G 25/4

ML





JMC FILE COL

REPORT DOCUMENTATION PAGE

4

AD-A191 391

		1d RESTRICTIVE MARKINGS. NONE	
		3 DISTRIBUTION/AVAILABILITY OF REPORT UNLIMITED	
4 PERFORMING ORGANIZATION REPORT NUMBER(S) 2		5. MONITORING ORGANIZATION REPORT NUMBER(S) 2	
6a. NAME OF PERFORMING ORGANIZATION The Trustees of Columbia University in the City of N.Y.		6d. OFFICE SYMBOL (If applicable)	
6c. ADDRESS (City, State, and ZIP Code) Office of Projects and Grants Box 20, Low Memorial Library Columbia University New York, New York 10027		7a. NAME OF MONITORING ORGANIZATION Office of Naval Research	
7b. ADDRESS (City, State, and ZIP Code) 33 Third Avenue Lower Level New York, New York 10003			
8a. NAME OF FUNDING/SPONSORING ORGANIZATION Office of Naval Research		8b. OFFICE SYMBOL (If applicable) N62927	
9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-85-K-0371			
10. SOURCE OF FUNDING NUMBERS			
PROGRAM ELEMENT NO.		PROJECT NO.	TASK NO.
800 North Quincy Street Arlington, Virginia 22217-5000			WORK UNIT ACCESSION NO.
11. TITLE (Include Security Classification) Problems in Integrated Networks			
12. PERSONAL AUTHOR(S) Professor Mischa Schwartz, Thomas E. Stern, Aurel A. Lazar			
13a. TYPE OF REPORT FINAL		13b. TIME COVERED FROM 4/1/85 TO 12/31/87	
14. DATE OF REPORT (Year, Month, Day) 88-02-15		15. PAGE COUNT 9	
16. SUPPLEMENTARY NOTATION The views, opinions and/or findings contained in this report are those of the authors and should not be construed as an official Department of the Navy position, policy or decision unless so designated by other documentation.			
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	
19. ABSTRACT (Continue on reverse if necessary and identify by block number) This final report covers work carried out under ONR Contract N00014-85-K-0371. The basic problems addressed concerned transmitting mixtures of traffic of disparate types over a variety of communication networks. The research fell into three categories: 1. Hybrid Multiplexing 2. Packet voice 3. Optimal flow control			
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION	
22a. NAME OF RESPONSIBLE INDIVIDUAL Prof. M. Schwartz, T.E. Stern, A.A. Lazar		22b. TELEPHONE (Include Area Code) 212-280-3125	
		22c. OFFICE SYMBOL	

DD FORM 1473, 84 MAR

83 APR edition may be used until exhausted.
All other editions are obsolete.

SECURITY CLASSIFICATION OF THIS PAGE

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution UnlimitedDTIC
SELECTE
FEB 24 1988
S H D

OFFICE OF NAVAL RESEARCH
FINAL REPORT

1 April 1985 through 31 December 1987

for

Contract NOOO14-85-K-0371

Problems in Integrated Networks

Principal Investigators:

Prof. Mischa Schwartz

Prof. Thomas E. Stern

Prof. Aurel A. Lazar

Department of Electrical Engineering
Columbia University
New York, NY 10027

Reproduction in whole, or in part, is permitted for any purpose of the United States Government.

This document has been approved for public release and sale; its distribution is unlimited.

ONR Final Report

This report consists of a research summary (Part I) followed by an index of publications (Part II).

I. Research Summary

The research under this grant falls into three categories, hybrid multiplexing, packet voice and optimal flow control, each of which is summarized below.

A. Hybrid Multiplexer

1. Overview

We have undertaken the design and implementation of a communications multiplexer that can be used to study certain problems that arise in conceptualized future telecommunications networks in which many different types of communications traffic share network transmission facilities. Whereas in the past, different types of traffic (e.g., voice, data, and video) were placed on separate transmission facilities, the trend that has been developing in the last few years is for common networks to be specified. Specifically, many standards have been developed which fall under the well recognized heading Integrated Services Digital Network (ISDN). Much of early ISDN to be introduced will perhaps maintain segregation between disparate traffic types; in which case classic traffic analysis techniques can predict the performance characteristics of the system. The present project focuses on a design that would come into commercial being as networks evolve further and depart more from current notions.

The multiplexers of interest combine the traffic from a number of sources onto a single high information rate conduit; from here the traffic enters the network. Regardless of original form, all the sources are encoded into digital data streams prior to multiplexing. The bit rate of the multiplexed signal will be a number larger than the average bit rates of the incoming sources. The sources may be of a packet-switched or circuit-switched nature; the multiplexer output is a string of frames of fixed time period in which data circuits precede data packets. The test bed, including all ancillary equipment such as data and traffic generators, is currently developed only far enough to study issues related to circuit-switched traffic, but can be expanded to study problems that arise in the hybrid switching environment.

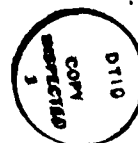
The multiplexer project seeks to blend knowledge from the study of teletraffic theory with the implementation techniques required to build a modern software controlled hardware system. In this way, performance evaluation theory can be verified.

The theoretical problem investigated is access control of the different circuit switched users vying for the bandwidth of the multiplexer. A user is characterized by his traffic load, how much bandwidth he uses, the length of his calls (holding time), and the rate at which he generates calls. Performance of the system may be measured by how much traffic is delivered to the network (throughput), or the probability that a user's requests for bandwidth are honored. With individual users having different characteristics, and competing for bandwidth:

- a) Different users may have varying quality with regards to blocking.
- b) The system throughput may be improved by restricting access of certain users.

2. Hardware/Software Description

The project is physically realized with several pieces of hardware and associated software. On the highest level the equipment can be split into the multiplexer per se, and the traffic generator. These are both constructed on Intel Multibus circuit boards, as shown in figure 1. Communications between the traffic generator and the multiplexer occur over the multibus.



<input checked="checked" type="checkbox"/>	
<input type="checkbox"/>	
<input type="checkbox"/>	
Availability Codes	
Avail and/or	
Dist	Special
A-1	

The multiplexer consists of several subsystems. On one circuit board is the hardware necessary to take several buffered sources of traffic and multiplex them onto a time division multiplex (TDM) frame. This is known as the frame generator. A single-board computer (Heurikon HK-68B) to drive the hardware and to make access control decisions also resides in the card cage and communicates with the frame generator over the multibus. It is known as the access controller.

Additional circuitry in the form of buffers is needed to interface the frame generator to real data sources. This circuitry has been designed for three separate cases: the ubiquitous RS-232 asynchronous data terminal, with a rate of between 300 bps and 19200 bps, the less common synchronous terminal, and analog voice, which would be digitized at the multiplexer. The sources must be buffered for a time equal to twice the TDM frame length. The buffers would reside on additional multibus circuit boards, although communication with the frame generator would occur on a local bus separate from the multibus.

The traffic generator consists of one single-board computer, identical to that used for the multiplexer. The traffic generator signals the multiplexer via multibus interrupts to indicate changes in on-hook/off-hook status of the emulated users.

Both the traffic generator and the multiplexer have a real time operating system known as MTOS (Multi-Tasking Operating System). Control of the multiplexer and the traffic generator is afforded from consoles as shown in figure 1.

Finally, links to a SUN workstation allow programs to be loaded into the traffic generator and multiplexer, and allow data to be collected from the multiplexer for offline analysis on the SUN workstation.

The MTOS operating system provides control over the system in two notable ways. One is the provision for a system console and links to a host computer (SUN workstation). This system control facet of operating system is used by both the traffic generator and the multiplexer. The other feature of note is the ability to perform different tasks simultaneously. Whereas both multiplexer and traffic generator are written as multitasking programs, the tasks themselves are specific to the application.

The traffic generator emulates up to sixteen circuit switched users. Each user has its own task. The task receives the parameters for the channel (for example, average wait and hold times) from the system operator at initialization time.

The task calculates a wait or hold time based upon an operator specified random distribution. The operating system pause function is used to wait the appropriate interval, and then an interrupt is sent to the multiplexer indicating ON/HOOK or OFF/HOOK. The task then terminates.

The multiplexer will interrupt the traffic generator to reflect the disposition of the traffic generator's request. The traffic generator's interrupt handler will re-start the appropriate task to generate the needed ON/HOOK or OFF/HOOK request after a newly calculated time.

The multiplexer software is somewhat more complicated. When a request arrives from the traffic generator, it is put on the appropriate list for disposition; either an ACCESS list or a DELETE list corresponding to OFF/HOOK or ON/HOOK respectively.

The bulk of the multiplexer's work is done in the access control task. This keeps track of how many users of what type are currently using the frame. It checks the DELETE list and then the ACCESS list to see if any work is to be performed. If users are to be added or deleted from the frame, the frame generation hardware must be manipulated, and the proper control messages must be sent over the TDM output. Although a user requesting deletion is always granted that, in the case of an access request, some algorithm must be applied to determine whether access is allowable. The algorithm is entered by the system operator when the system is initialized. Access control keeps track of the frequency with which the requests are accepted or denied so that probability of loss information can be sent to UNIX periodically (the SEND task sends the information).

Another task runs periodically every 1/4 second and keeps a log of how many users are occupying the frame. Every fifteen minutes the SEND task sends this log to the SUN workstation.

3. Analytical Work

The analytical work has focused on the problem of maximizing the throughput of a circuit switched facility when the facility is shared by two types of users having different bandwidths. The multiplexer is designed so that complete sharing of the band is possible among the different classes of users: there are no physical restrictions limiting certain types of users to certain parts of the band. Given that the resource can be shared in this manner, the question arises as to the performance of the system with complete sharing.

A restricted access policy would limit the number of users allowed on the band even when room exists for more. A result was found for the case in which the bandwidths of the two user classes form an integer ratio. To achieve maximum throughput one should limit the number of the narrowband users but not restrict the number of wideband users. Also, there are only particular values that this limit should take, based upon the ratio of the bandwidths and the total bandwidth available. The number to which the narrowband user is limited depends upon the traffic parameters of the two user classes. In fact, the limit may be large enough to "disappear"; complete sharing is optimal with some traffic parameters. Although we have determined the nature of the policy, a simple relationship between traffic parameters and choice of limit value does not exist. The problem can be solved numerically on a case by case basis.

4. Results of Emulation

A number of experiments have been run with emulation to determine whether analytical results can be duplicated on real hardware.

One set of experiments involved measuring the distributions of users on the system given the Engset traffic situation with one class of user. There are m users trying to access n channels. All users have the same waiting time and holding time distributions. In the initial experiments, wait and hold times were exponentially distributed. We measured both the fraction of time that i users are on the frame, $0 \leq i \leq m$, and the probability of loss (the fraction of calls that are blocked). Following the case in which wait and hold times are exponentially distributed, some other distributions were substituted for one or both of the wait and hold times. Other distributions tried were uniform distributions and fixed (discrete) distributions. In all cases the probabilities corresponding to the Engset distribution were obtained, except when both wait and hold times were fixed distributions. This is as would be expected.

The next group of experiments considered two classes of circuit-switched users with different bandwidths vying for the resource. The goal was to see whether the restricted access policy predicted by the theory mentioned above would yield maximum throughput. The experiments showed that when the bandwidths of the two classes were a large ratio (the experiment was performed with 64 kbps users and 9.6 kbps users), restricted access could yield a large gain in throughput over complete sharing. We also showed that with a small ratio of bandwidths (2:1), it was not possible to measure a large improvement in throughput by using restricted access policies. It became apparent also that the policy which provides the best performance in terms of throughput may cause poor performance in terms of probability of loss to one of the users. It is conceivable to have maximum throughput when the narrowband user is limited so that only wideband traffic is admitted to the network. We are looking at graphical techniques to balance gains in throughputs with excessive disparities in loss performance between the user classes.

We have not been able to come up with a simple form of optimal access control strategy when there are more than two classes of users. We have attempted several ad hoc strategies for cases in which there are three or four different classes. An algorithm for the case of three types is to order the bandwidths from smallest to largest and then determine between which two groups is the largest difference in bandwidth. Restrict the group or groups below the largest gap. For the case of more than three users, a software algorithm is being developed to allow convenient and flexible partitioning of the user types into classes for different restricted access policies.

B. Packet Voice

Transmission of voice in packetized form offers a number of advantages over the conventional circuit-switched mode of transmission. These range from the economies of sharing common hardware, software and transmission facilities to the vast potential for new multimedia services involving the combined use of voice, images and data. The packet voice research completed under this ONI grant has both an analytical and experimental component. The analytical work has focused on the conception and evaluation of new multiplexing techniques for achieving improved quality of transmission while at the same time increasing network utilization. The experimental work has involved the implementation of a packet voice transmission system on our Integrated Local Area Network Testbed.

In the analytical work two projects have been completed involving the performance of systems using embedded coding.

1. Performance Analysis of Priority-oriented Packet Voice Systems.

We have proposed a priority-oriented system to improve voice transmission quality in packet switched networks. The system uses an imbedded coding technique in which two streams of information are generated, one more significant than the other, with the more significant information assigned a higher priority. Thus, it is the low priority packets that incur the highest queueing delay, and eventually the most losses due to excessive delays and buffer overflow. Such a system can provide "graceful degradation" of speech quality during periods of congestion by expediting the transmission of the more significant voice information at the expense of the less significant information. We developed a queueing model as an analytical tool for performance evaluation. Both arrival and service rates of the queue are modulated by a Markov chain (a birth-death process). The equilibrium queue length distribution was derived, and the mean waiting times and queue lengths of the different priority queues were shown to be related by a conservation law. The analysis, reported in [1] indicated that the system can indeed achieve significant improvements in performance based on the criteria of mean queue length and delay in the high priority stream. To align our theoretical results with subjective voice quality criteria we also constructed a real time packet voice simulator (implemented on a Masscomp computer), which simulates all of the operations in a packet voice system, including impairments such as queueing delays and packet loss. The resultant voice quality can be monitored by a listener for performance comparisons. Our analytical results were confirmed by informal subjective listening tests using the simulation facility.

2. Congestion Control by Selective Packet Discarding.

In this work we have developed a family of congestion control schemes for reducing time delays as well as memory requirements in packet voice systems. They are all based on selectively discarding packets whose loss will produce the least degradation in quality of the reconstructed voice signal. As in the priority-oriented schemes, it is assumed that the voice information is transmitted using two classes of packets, with Class 1 containing the more significant information, and Class 2 the less significant information. The Class 2 packets are selectively discarded when the onset of congestion occurs. Queueing models similar to those used in the priority-oriented systems we used to evaluate performance in terms of packet delay and loss. A variety of control procedures were examined and compared. They involved different methods of classifying the encoded voice, and different packet discarding rules. Packets were classified as containing the most or least significant bits in PCM samples, as containing even or odd samples and in terms of speech energy detection thresholds. The packet discarding rules were based on speaker activity thresholds or buffer content thresholds, and included the use of a "resume level". The analysis, reported in [2], shows that this type of congestion control produces substantially improved performance over uncontrolled systems, based on criteria of mean waiting time and fractional packet loss. Again, the quantitative work was validated by subjective testing using our simulator.

3. Implementation of Packet Voice on an Integrated LAN.

In this work we have developed a packet telephone system, that uses an Integrated LAN for communications. Plain Ordinary Telephone Service (POTS) is supported by the system, meaning that stations on the LAN may call each other, and have a conversation. Flexible and tested algorithms were implemented for speech activity detection and speech reconstruction. Several algorithms for speech compression were tested, and rejected. An acceptable level for background noise, and for mouth to ear delay was established.

System architecture is as follows. A Voice Interface Unit (VIU) links to the telephone instrument, and provides A/D and D/A conversion as well as separation of signaling information. A Terminal Interface unit (TI) provides access to the LAN used, the MAGNET Integrated Local Area Network. The Voice Controller subsystem (VC) is a satellite processor responsible for packetizing voice samples read from the VIU, and transmitting them over the LAN via the TI. For the opposite direction of data flow, it is responsible for receiving packets and reconstructing the byte stream. In addition, this controller is responsible for call set up and take down. It is implemented as a single-board computer, with the real-time executive MTOS, stored as firmware, running to support voice service responsibilities. This executive provides an environment in which to execute tasks, each of which assumes the work burden for a portion of the packet telephone system. The executive also enables the tasks to communicate with each other, and, via drivers, to the VIU and the TI.

The software system is illustrated in figure 2. The tasks and drivers have been split along three functional planes: the M-plane for resource monitoring, the C-plane for connection set up and take down, and the U-plane for user information transmission. The U-plane sends and receives packets of voice samples, and reconstructs the voice byte stream. Here the Voice Controller single-board computer receives voice samples from the VIU as bytes read directly into its memory by an on board DMA chip. When a full buffer is accumulated, the Telephone Mouthpiece Driver activates the Packetization Task, and presents it the buffer. This task tests the buffer using a speech activity detection algorithm, and if the buffer is deemed to contain active speech, sends it in a packet to the TI Output Driver, which transmits the packet over the LAN. The Voice Controller presents voice samples to the VIU by having buffers ready for the DMA, which extracts the samples directly from memory. Byte stream reconstruction is accomplished by the Telephone Earpiece Driver which orders the buffers it receives and inserts buffers of silence where necessary. The received buffers are readied by the Depacketization Task, which receives packets from the TI Input Driver. The C-plane establishes a connection between stations, and starts and stops the U-plane. The protocols are connection oriented, and provide such primitives as connection request, connection accept or reject, disconnection request, and disconnection accept. Some of the layer boundaries, as defined by the OSI communications protocol standards, are indicated on the figure. The M-plane generates simple reports on the number of packets sent and received, for display by another station.

C. Optimal Flow Control

Under this and the previous ONR grant, we have been developing a unified treatment of the problem of optimal routing and flow control in computer communications networks. In recent work, a Markovian queueing model has been used to analyze decentralized flow control mechanisms in networks with multiple controllers. If network performance is to be optimized, the optimal decentralized flow control problem becomes a team decision problem. If each user's performance is to be optimized, the problem becomes a multi-person game. The conditional estimates of the total arrival and departure rates (respectively, user service rate) are sufficient statistics for the network (respectively, user) optimization problem. Using linear programming, the network optimal control and the Nash equilibrium solution for the game theoretic formulation are a set of window-type mechanisms. We have characterized the class of decentralized flow control problems with Nash equilibrium solutions.

In related work extending our earlier results on routing and flow control in a network of parallel processors, we have solved a general load sharing problem in Markovian queueing networks. This work was reported in [3] and [4].

II. Index of Publications

- [1] N. Yin, T.E. Stern and S.Q. Li, "Performance Analysis of a Priority-Oriented Packet Voice System," INFOCOM '87, San Francisco, Mar. 1987.
- [2] N. Yin, S.Q. Li and T.E. Stern, "Congestion Control for Packet Voice by Selective Packet Discarding," Globecom '87, Tokyo, Japan, Nov. 1987.
- [3] A. Bovopoulos and A. Lazar, "Decentralized Algorithms for Optimal Flow Control," Proc. Allerton Conf. on Communications, Control and Computers, Urbana Ill., Oct. 1987.
- [4] A. Bovopoulos and A. Lazar, "Optimal Load Balancing for Markovian Queueing Networks," Proc. Midwest Symposium on Circuits and Systems, Syracuse, N.Y., Aug. 1987.

- [5] B. Kraimeche, M. Schwartz, "Bandwidth Allocation Strategies in Wide-Band Integrated Networks," IEEE J. on Sel. Areas in Comm., Sept. 1986.
- [6] B. Kraimeche, M. Schwartz, "Analysis of Traffic Control Strategies in Integrated Services Networks," IEEE Trans. on Comm., Oct. 1985.
- [7] S. Ganguly, "Analysis and Performance Evaluation of a Packet Voice System," (submitted for publication).
- [8] A.A. Lazar, and M.T. Hsiao, "Network and User Optimal Flow Control with Decentralized Information," (submitted for publication).
- [9] M. Luhanga and T.E. Stern, "Analytical Modeling of Small Packet Voice Systems," Int. J. Elect. Enging. Educ. Vol. 22, pp. 339-344.

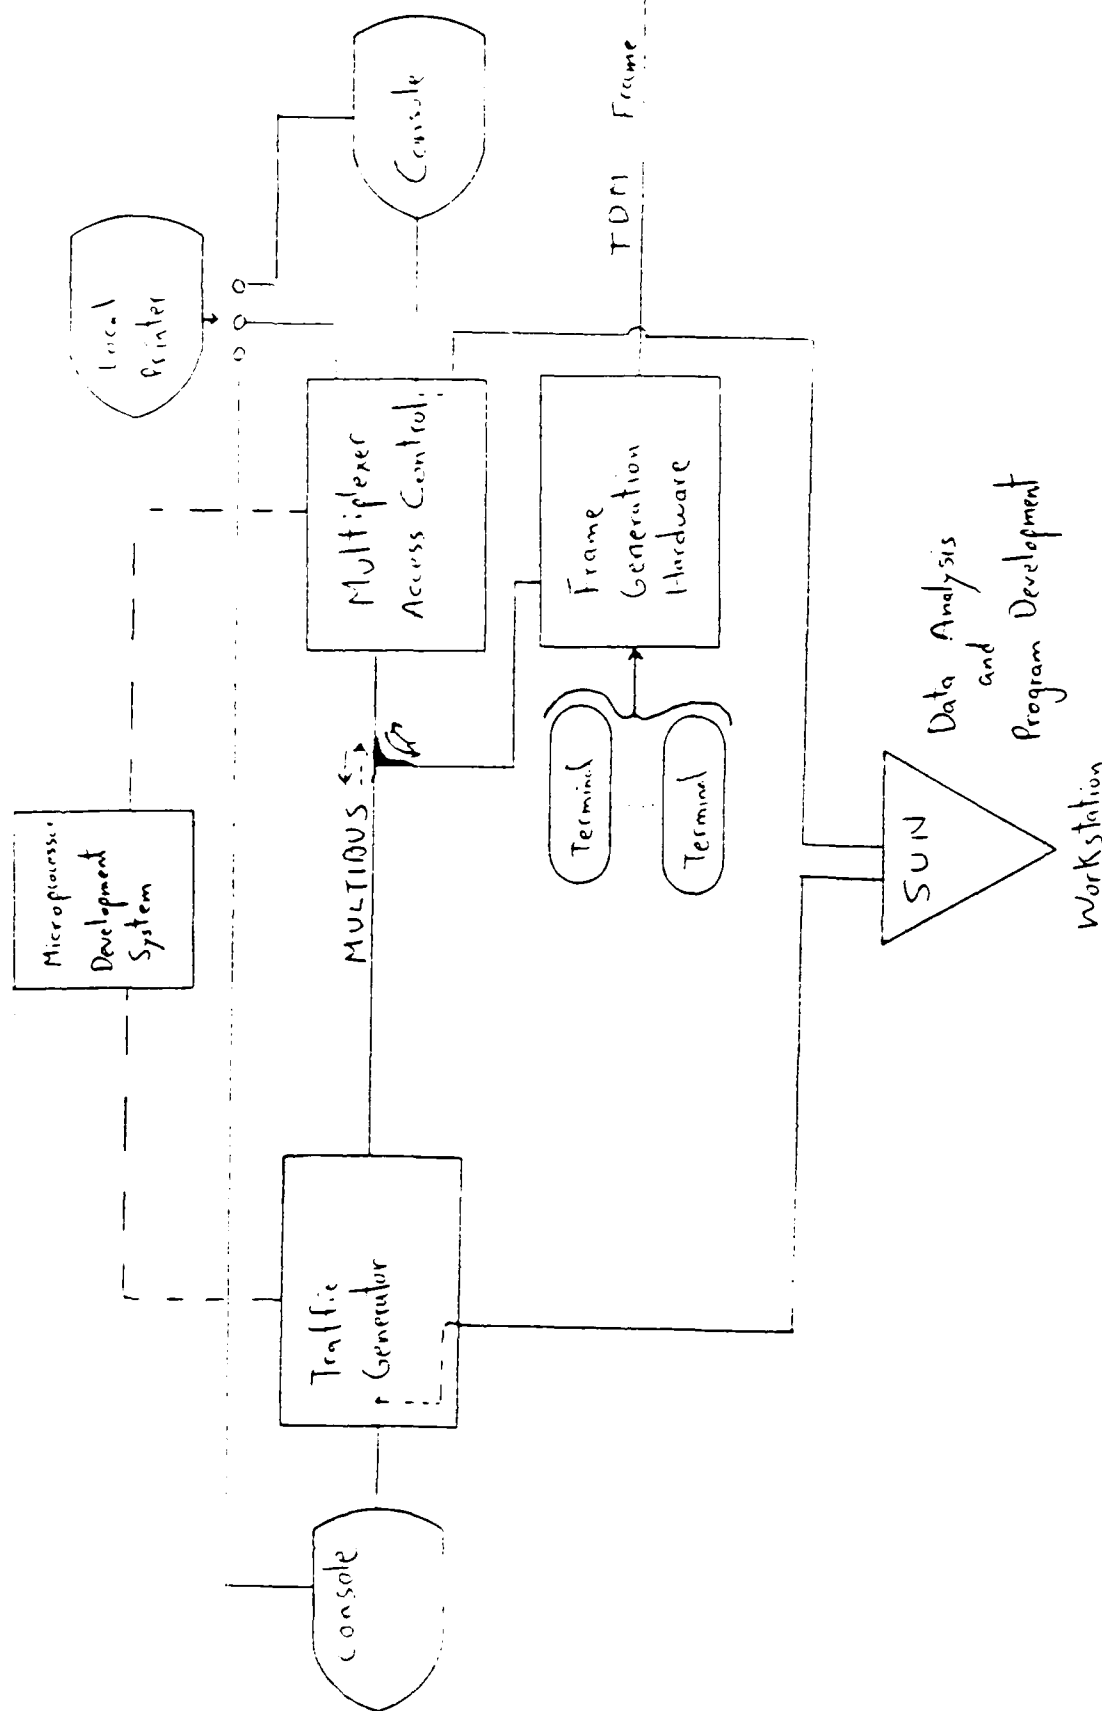


Figure 1

Overall Physical Configuration of
Hybrid Multiplex Project

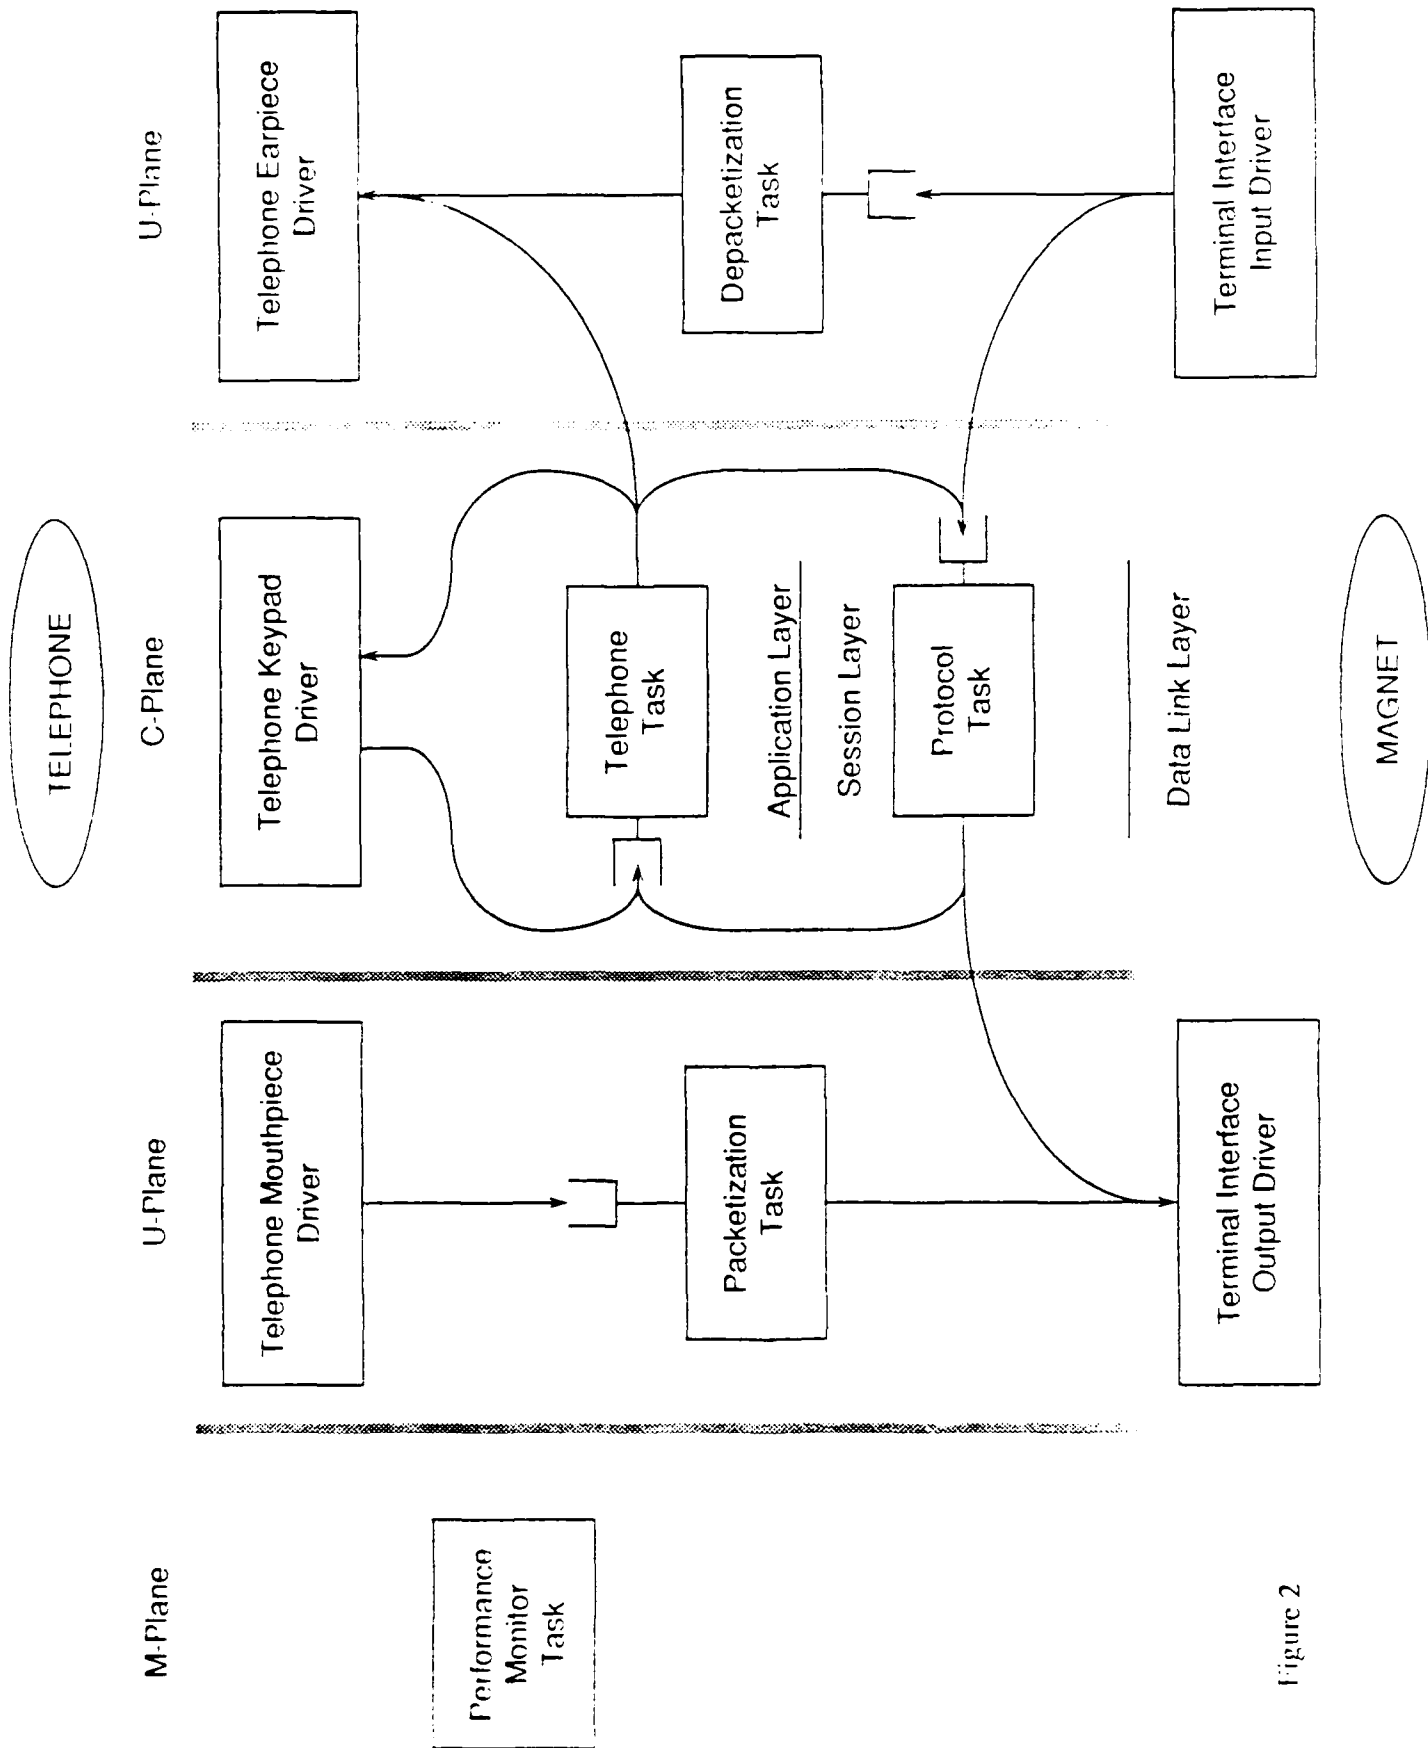


Figure 2

END

DATE

FILMED

5-88
DTIC